



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁶ : G06F 17/60</p>	A1	<p>(11) International Publication Number: WO 99/63466</p> <p>(43) International Publication Date: 9 December 1999 (09.12.99)</p>						
<div style="display: flex; justify-content: space-between;"> <div style="width: 48%;"> <p>(21) International Application Number: PCT/US99/12347</p> <p>(22) International Filing Date: 3 June 1999 (03.06.99)</p> <p>(30) Priority Data:</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 30%;">09/092,348</td> <td style="width: 30%;">5 June 1998 (05.06.98)</td> <td style="width: 40%;">US</td> </tr> <tr> <td>09/156,265</td> <td>18 September 1998 (18.09.98)</td> <td>US</td> </tr> </table> <p>(71) Applicant: i2 TECHNOLOGIES, INC. [US/US]; Suite 1600, 909 East Las Colinas Boulevard, Irving, TX 75039 (US);</p> <p>(72) Inventors: NOTANI, Ranjit, N.; 1218 Hidden Ridge Drive #2070, Irving, TX 75039 (US). PARASNIS, Abhay, V.; 8804 Saddlehorn Drive #208, Irving, TX 75063 (US). WHIPPLE, Mark, B.; 110 North Clinton Avenue, Dallas, TX 75208 (US).</p> <p>(74) Agent: MEEK, Kevin, J.; Baker & Botts, L.L.P., 2001 Ross Avenue, Dallas, TX 75201-2980 (US).</p> </div> <div style="width: 48%;"> <p>(81) Designated States: AE, AL, AM, AT, AT (Utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published</p> <p><i>With international search report.</i></p> <p><i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p> </div> </div>			09/092,348	5 June 1998 (05.06.98)	US	09/156,265	18 September 1998 (18.09.98)	US
09/092,348	5 June 1998 (05.06.98)	US						
09/156,265	18 September 1998 (18.09.98)	US						
<p>(54) Title: SYSTEM AND METHOD FOR IMPLEMENTING OBJECT WORKSPACE AGENTS IN A DECISION SUPPORT ENVIRONMENT</p>								
<p>(57) Abstract</p> <p>A computer system for remotely accessing data in a multi-enterprise collaboration comprises a workspace associated with a first enterprise having a plurality of stored objects. The computer system further comprises a network node associated with a second enterprise, the network node being in communication with the workspace across the network. The computer system further comprises an agent generated at the network node, the agent operable to access the workspace via the network, the agent further operable to manipulate at least one of the plurality of stored objects within the workspace to perform a collaboration activity.</p>								

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

SYSTEM AND METHOD FOR IMPLEMENTING OBJECT WORKSPACE
AGENTS IN A DECISION SUPPORT ENVIRONMENT

TECHNICAL FIELD OF THE INVENTION

This invention relates in general to the field of supply chain, enterprise and site planning and, more particularly, to a system and method for implementing
5 object workspace agents in a decision support environment.

BACKGROUND OF THE INVENTION

Supply chain, enterprise and site planning
10 applications and environments are widely used by manufacturing entities for decision support and to help manage operations. Decision support environments for supply chain, enterprise, and site planning have evolved from single-domain, monolithic environments to multi-
15 domain, monolithic environments. Conventional planning software applications are available in a wide range of products offered by various companies. These decision support tools allow entities to more efficiently manage complex manufacturing operations. However, supply chains
20 are generally characterized by multiple, distributed and heterogenous planning environments. Thus, there are limits to the effectiveness of conventional environments when applied to the problem of supply chain planning due to monolithic application architectures. Further, these
25 problems are exacerbated when there is no one "owner" of the entire supply chain.

It is desirable for the next evolutionary step for planning environments to establish a multi-domain, heterogenous architecture that supports products spanning multiple domains, as well as spanning multiple engines and products. The integration of the various planning environments into a seamless solution can enable inter-domain and inter-enterprise supply chain planning. Further, an important function provided by some planning applications is the optimization of the subject environment rather than simply tracking transactions. In particular, the RHYTHM family of products available from I2 TECHNOLOGIES provide optimization functionality. However, with respect to planning at the enterprise or supply chain level, many conventional applications, such as those available from SAP, use enterprise resource planning (ERP) engines and do not provide optimization.

The success or failure of an enterprise can depend to a large extent on the quality of decision making within the enterprise. Thus, decision support software, such as I2 TECHNOLOGIES' RHYTHM family of products, that support optimal decision making within enterprises can be particularly important to the success of the enterprise. In general, optimal decisions are relative to the domain of the decision support where the domain is the extent of the "world" considered in arriving at the decision. For example, the decision being made may be how much of a given item a factory should produce during a given time period. The "optimal" answer depends on the domain of the decision. The domain may be, for example, just the factory itself, the supply chain that contains the factory, the entire enterprise, or the multi-enterprise supply chain. (The latter two can be considered to be larger domains or

multiple domains.) Typically, the larger the domain of the decision support, the more optimal the decision will be. Consequently, it is desirable for decision support software to cover ever larger domains in the decision making process. Yet, this broadening of coverage can create significant problems.

One such problem involves the limitations of present decision support software that do not allow remote entities in the domain or multi-domain environment to access a local object workspace. Such software does not allow such remote entities to easily execute administrative tasks, optimization routines or other forms of data and object manipulation in a local workspace.

SUMMARY OF THE INVENTION

In accordance with the present invention, a system and method for implementing object workspace agents in a decision support environment are disclosed that provide advantages over conventional supply chain, enterprise and site planning environments.

According to one aspect of the present invention, a computer system for remotely accessing data in a multi-enterprise collaboration comprises a workspace associated with a first enterprise having a plurality of stored objects. The computer system further comprises a network node associated with a second enterprise, the network node being in communication with the workspace across the network. The computer system further comprises an agent generated at the network node, the agent operable to access the workspace via the network, the agent further operable to manipulate at least one of the plurality of stored

objects within the workspace to perform a collaboration activity.

A technical advantage of the present invention is the ability to access a remote workspace with an agent to satisfy programmed goals of the agent in that remote workspace. A further technical advantage of the present invention is the ability of the agent to execute tasks in a first enterprise, node or platform despite being initiated at a second enterprise, node or platform with which the agent has lost contact. A further advantage of the invention is that the agent is able to communicate results of executed tasks performed at a destination node back to an originating node. Another advantage is that the agent is operable to execute or operate on other agents which are residing in a workspace. Additional technical advantages should be readily apparent to one skilled in the art from the following figures, descriptions, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the present invention and advantages thereof may be acquired by referring to the following description taken in conjunction with the accompanying drawings, in which like reference numbers indicate like features, and wherein:

FIGURE 1 is a diagram of one embodiment of a computer implemented architecture that can support enterprise collaboration;

FIGURE 2 is a diagram of one embodiment of components of a global collaboration framework;

FIGURE 3 is a diagram of the global collaboration framework of FIGURE 2 where certain software elements that make up particular modules are highlighted;

FIGURE 4 is a block diagram of one embodiment of a system allowing collaboration within and between enterprises for optimal decision making;

5 FIGURE 5 is a block diagram of one embodiment of the use of a global collaboration workspace;

FIGURE 6 is a diagram of one embodiment of a lifecycle for a collaboration;

10 FIGURE 7 is a diagram of situations where common software is present on both sides of a relationship and where it is not;

FIGURE 8 is a block diagram of one embodiment of a security configuration for a hub-to-spoke and hub-to-web case;

15 FIGURE 9 is a block diagram of one embodiment of a security configuration for a hub-to-hub case;

FIGURE 10 is a diagram of one embodiment of designing an inter-enterprise workflow that includes parameterization over groups;

20 FIGURE 11 is a diagram of one embodiment of managing change by modifying a design of a workflow;

FIGURE 12 is a diagram of one embodiment of integration of a workflow with the outside world;

FIGURE 13 is a diagram of one embodiment of a data flow running in a single activity;

25 FIGURE 14 is a diagram of one embodiment of a data flow split across multiple activities;

FIGURE 15 is a block diagram of one embodiment of an common data model based transformation model;

30 FIGURE 16 is a diagram of one embodiment of a direct transformation;

FIGURE 17 is a diagram of one embodiment of different access and transformation levels;

FIGURE 18 is a diagram of one embodiment of substituting a hub engine for a spoke engine within a collaboration;

5 FIGURE 19 is a block diagram on one embodiment of a computer system using a workspace configured according to the teachings of the present invention;

FIGURE 20 is a diagram of one embodiment of the workspace of FIGURE 19 configured according to the teachings of the present invention;

10 FIGURE 21 is a block diagram of a computer system using an agent to access an object workspace according to the teachings of the present invention; and

15 FIGURE 22 is a flowchart illustrating one embodiment of a method of using the agent of FIGURE 21 to access a workspace.

DETAILED DESCRIPTION OF THE INVENTION

Improvement of decision support processes involves expansion to provide enterprise level and multi-enterprise level decision support for optimal decision making. Technologically and conceptually, providing enterprise-level and multi-enterprise level decision support differs from providing factory-level and supply-chain-level decision support. The reasons for this can be that, in multi-domain situations (such as business units within an enterprise or multiple enterprises), the different domains often operate different decision support software. Also, in multi-domain situations, one domain generally can not coerce another domain into making a particular decision. In other words, optimal decision support in this environment often needs to be performed in a negotiated, as opposed to coercive, environment.

20
25
30

Providing decision support in multi-domain situations can be accomplished by pursuing a collaborative approach to decision support rather than a coercive one. Various communication and distributed processing technologies can be used to implement such an environment, including the Internet, the Web, JAVA, XML, CORBA, etc., which help make large scale collaborative decision making feasible. Products will soon be available from I2 TECHNOLOGIES that enable a collaborative approach to decision support, including RHYTHM-GLOBAL COLLABORATION MANAGER (GCM) and RHYTHM-GLOBAL COLLABORATION DESIGNER (GCD).

Collaboration System and Process Components

FIGURE 1 is a diagram of one embodiment of a computer implemented architecture that can support enterprise collaboration. As shown, a global decision support architecture can be built upon underlying link, vision, global messaging and data warehouse components. Collaboration can then involve a global collaboration designer (GCD) and a global collaboration manager (GCM) supported by the decision support architecture. The global collaboration designer can be used to design and instantiate collaborations, and the global collaboration manager can be used to run the collaborations. In this scheme, collaborations can be referred to as modules and can be versioned.

FIGURE 2 is a diagram of one embodiment of components of a global collaboration framework. As shown, the framework can allow an hub enterprise 2 to collaborate with a spoke enterprise 4 and a web enterprise 6. Hub enterprise 2 and spoke enterprise 4 each include a global collaboration manager 8. Global collaboration managers 8

are coupled to and communicate with respective internal global collaboration workspaces 10. An external global collaboration workspace 12 provides a means for sharing data between hub enterprise 2, spoke enterprise 4 and web enterprise 6. Hub enterprise 2 can also collaborate through an electronic data interchange (EDI) processor 14 with a value added network (VAN). Further, hub enterprise 2 can communicate and collaborate with other hub enterprises using a global message bus 15.

In operation, the primary controller of the collaboration can be the GCM engine 8 of hub enterprise 2. The hub-to-hub relationship can be facilitated by the global message bus 15, and the hub-to-spoke and hub-to-web relationships can be facilitated by external global collaboration workspace (GCW) 12. As shown, a hub enterprise 2 can generally have an internal GCW 10 and an external GCW 12. Internal GCW 10 can be used to share and exchange data with internal user interfaces as well as EDI processor 14. External GCW 12 can be used to share and exchange data with spoke enterprises 4 and web enterprises.

For security, external GCW 12 can be installed in a DMZ or outside a corporate firewall of hub enterprise 2. This way no direct connections need to be made from the outside into the protected corporate network of hub enterprise 2. External GCW can accept, for example, IIOP, HTTP and HTTPS connections. In particular, the latter two connections are useful for bridging existing firewall configurations. In this manner, no firewall configuration is needed on either the client (spoke node or web node) or server (hub node) side which can make the solution more quickly deployable.

FIGURE 3 is a diagram of the global collaboration framework of FIGURE 2 where certain software elements that make up particular modules are highlighted. As can be seen, software for the global collaboration manager module can be present in the following places: in the hub engine 8, in the spoke engine 8, in the hub-user user interface (UI), in the spoke-user UI and in the web-node UI. Additionally, the module can communicate with native applications 17 on the hub enterprise 2 and spoke enterprise 4. Communications with native applications 17 can be either synchronous (dot line) or asynchronous (solid line). Asynchronous communication with native applications 17 can be facilitated by the internal GCW's 10, as shown. Further, a global series database (GSDB) can be present on the hub enterprise 2 side.

FIGURE 4 is a block diagram of one embodiment of a system, indicated generally at 16, allowing collaboration within and between enterprises for optimal decision making. As shown, system 16 includes a hub node 18 which can be a process within a hub engine executing on a computer system. Hub node 18 is coupled to and communicates with a spoke node 20 which also can be a process within a hub engine executing on a computer system. As shown, spoke node 20 can be outside an enterprise boundary 22 of hub node 18. Hub node 18 is also coupled to and communicates with a plurality of spoke nodes 24 which can be processes within a spoke engine executing on one or more computer systems. Hub node 18 can further be coupled to and communicate with a plurality of web nodes 26 which can be processes within a web browser executing on a computer system. In addition, hub node 18 is coupled to and communicates with an EDI

(Electronic Data Interchange) proxy 28 which can provide a gateway to EDI systems.

Hub engines and spoke engines, together with a global collaboration workspace, can be the primary entities of a global collaboration manager. In this environment, a hub engine is the primary controller of the collaboration. The hub engine can coordinate both global collaborations as well as local collaborations. Global collaborations are those that span hub nodes 18, spoke nodes 20 and 24 and web nodes 26. Local collaborations can run on any single role hub or spoke/spoke-group. These collaborations can be distributed, but stay within the confines of a single enterprise. Hub engines can also coordinate with hub-user interfaces (UI) as well as the VAN-EDI processor of an EDI proxy 28. In one embodiment, hub engines are multi-threaded engines that can simultaneously coordinate multiple collaborations as well as multiple versions of the same collaboration. Further, the hub engines can dynamically load and execute collaborations.

A spoke engine can also operate to initiate a collaboration. In this environment, unlike a hub engine, a spoke engine is not an independent entity. Instead a spoke engine can only coordinate a collaboration in conjunction with a hub engine. Furthermore, a spoke engine can not coordinate with other spoke engines or other web-nodes. Like a hub engine, a spoke engine can be multi-threaded and can simultaneously coordinate multiple collaborations as well as multiple versions of the same collaboration. Spoke engines can also dynamically load and execute collaborations.

FIGURE 5 is a block diagram of one embodiment of the use of a global collaboration workspace 30. In FIGURE 5,

global collaboration workspace 30 provides the primary entity used to share data/objects between the various entities in the collaboration. As shown, workspace 30 can interface with global collaboration managers (GCMs) 32, a
5 local system 34, a web server 36 and web interface 37 and native applications 38. In general, objects can be placed into global collaboration workspace 30 by one entity and retrieved by another entity. Retrieval can be achieved either by querying or by subscription. In this way, global
10 collaboration workspace 30 combines the attributes of a database as well as a message bus.

The global collaboration workspace can be organized as a hierarchy of slots which can be in-memory or persistent. Slots also can be queued or regular, and fine grained
15 permissibilities can be attached to each slot. The permissibilities can be assigned by-user-by-operation. The primary operations can be read, write, take, and subscribe.

In-memory slots hold their data in volatile memory. Writing and retrieval from in-memory slots is very fast but
20 subject to loss if the global collaboration workspace 30 goes down. When used with in-memory slots, the global collaboration workspace 30 can be considered a fast, secure, in-memory object database, with security and messaging capabilities. Persistent slots hold their data
25 in stable storage. Writing and retrieval from persistent slots is slower than for in-memory slots, but data is not lost if the global collaboration workspace 30 goes down.

The decision as to whether to use in-memory or persistent slots can depend on the application. Global
30 collaboration workspace 30 stores data in the form of objects and can store Java Objects, CORBA objects or arbitrary byte arrays. This, coupled with its in-memory

capabilities, makes global collaboration workspace 30 appropriate as a high-speed data sharing mechanism between other object-oriented in-memory engines such as I2 TECHNOLOGIES' SUPPLY CHAIN PLANNER and FACTORY PLANNER.

5 A global collaboration designer (GCD) provides a tool to allow collaboration designers to interactively design, instantiate and deploy collaborations to be run using the global collaboration manager. The output of the global collaboration designer is code that can be automatically
10 loaded and run by the global collaboration manager. The global collaboration designer can allow designers to create new collaborations, retrieve existing collaborations, and version collaborations. The global collaboration designer can also allow designers to design the hub and spoke
15 network for collaborations and design the events and messages of the collaboration. The global collaboration designer can integrate a standard object library and a standard component library for easy usage from within the global collaboration designer. The global collaboration
20 designer can be used to create sophisticated multi-enterprise workflows with synchronous, asynchronous, sub-workflow, and-splits, or-splits, synchronization-joins, heterocast-splits, heterocast-joins etc. Global workflows and local workflows can both be created. The global
25 collaboration designer can provide automatic verification of collaborations and automatic code generation, which code is run by the global collaboration manager. The generated code can be manually edited if desired. Further, the global collaboration designer can provide instantiation of
30 a collaboration including generation of security manager configurations and global collaboration workspace configurations.

FIGURE 6 is a diagram of one embodiment of a lifecycle for a collaboration. As shown, in step, a collaboration can be designed using the global collaboration designer. Then, in step 46, a collaboration can be instantiated using the global collaboration designer. The instantiated collaboration can then be deployed, in step 44, using the global collaboration designer and the global collaboration manager. After deployment, the collaboration can be run using the global collaboration manager in step 46. Subsequently, a new instance can be created or a new version of the collaboration can be created. To create a new instance, the flow returns to step 42. For a new version, the global collaboration designer can be used in step 48 to modify the collaboration.

The extension from single-domain decision support to multi-domain decision support can be complicated. In particular, the following discussion describes a number of challenges presented by multi-domain decision support and embodiments of how those challenges are addressed by the present system and process allowing collaboration within and between enterprises for optimal decision making.

Representational Heterogeneity

One problem with collaboration is bridging representational heterogeneity across enterprises. Before collaboration can successfully occur, the representational heterogeneity across enterprises needs to be bridged. Enterprises often represent the same data in different ways. These differences range from semantic differences, to technological differences, to differences in naming, etc. One obvious solution to bridging these differences is standardization. However, this immediately raises the

issue of what standard to agree upon. The present system and process avoid such a requirement.

It should be noted that there can be three relevant categories of standards that need to be addressed. These three categories are: format standards, transport standards and semantic standards. Format standards refer to the technological formats in which the data/objects are encoded. Examples include XML, Java Serial Streams, IIOP Serial Streams and EDI format. Transport standards are used to pass data around. These can include HTTP, IIOP, RMI, DCOM, FTP, Value Added Networks, Asynchronous Message Buses such as MQSeries, etc. Third, semantic standards are the way in which the semantic content of the data is described. Examples include EDI, I2 COMMON DATA MODEL (CDM).

By considering standards in this light, an understanding of the issues can emerge. A lot of the confusion today stems from the fact that many existing standards cover two or more of the categories above and that discussions of the various standards fail to categorize which category is being discussed. For example, EDI is primarily a semantic standard, but also typically implies a format standard (the EDI file format) and a transport (a Value Added Network). Once this is understood, it becomes clear that the EDI semantic standard can be separated from the other two. Hence, semantic EDI objects can be encoded in other formats such as Java Serial Streams and can be passed over other transport standards such as HTTP. Similarly, XML is primarily a format standard that can be used to encode various semantic standards. Efforts are underway to encode EDI in XML.

Several format standards can be supported by the present global collaboration manager, including XML, EDI format, Java Serial Streams (referred to as Java format and not to be confused with the Java Language or Java Platform) and IIOP Serial Streams. Of these, in one embodiment, the Java format is the primary format, and the rest are derived formats. Because the Java Format can contain the behavior to produce the other formats, it has been chosen as the primary format. XML, EDI and IIOP formats can be derived from the Java Format.

FIGURE 7 is a diagram of situations where common software from I2 TECHNOLOGIES' is present on both sides of a relationship and where it is not. As shown, for example, when RHYTHM GLOBAL COLLABORATION MANAGER is on both sides, nothing is to be gained by converting to an intermediate format. This would introduce needless inefficiency, and only data (not objects) would be exchangeable, limiting the range of applications. Hence when the same software is present on both sides, binary Java objects can be directly exchanged. On the other hand, for example, when RHYTHM GLOBAL COLLABORATION MANAGER is present only on one side, XML or EDI-formatted "objects" can be produced (outbound) and interpreted (inbound).

With respect to transport standards, the present global collaboration manager can support a variety of transport standards, including HTTP, IIOP, and Asynchronous Message Buses. More details are provided below with respect to Handling Multiple Relationship Types.

With respect to semantic standards, the present global collaboration manager can primarily support two semantic standards, EDI and RHYTHM-CDM. EDI can be supported because it is generally the most popular semantic standard.

However it suffers from the drawback (amongst others) of not providing deep coverage of the planning domain. The RHYTHM-CDM, on the other hand, provides deep coverage of the planning domain and will provide appropriate constructs for performing multi-enterprise decision support. Additionally, this format is supported by all of I2 TECHNOLOGIES' planning engines.

In general, one problem with public standards, such as EDI, is that they may not adequately cover the kinds of data/objects that enterprises would like to exchange. Further, waiting for standards bodies to standardize on a particular object may not be an option, and a supply chain will not have any particular competitive advantage by using public standards. For these and other reasons, the present global collaboration manager supports an alternative approach to standardization by supporting proprietary community standards. For example, using RHYTHM-GCD, a community of enterprises can devise a set of standards that are relevant to that community only. RHYTHM-GCM will support and enforce these proprietary community standards. RHYTHM-GCD also supports a library of building block objects that can be composed into proprietary community standards. Proprietary community standards have a number of advantages, including: they can be designed to exactly cover the kinds of data/objects that enterprises would like to exchange; only the relevant parties need to agree upon the particular standard, hence the process will be much quicker than waiting for a standards body; different standards can be developed for different categories of partners and, in the extreme case, a different standard for each partner; and standards that give the supply chain a competitive advantage over competitors can be developed.

Multiple Relationship Types

Another problem for allowing collaboration is handling multiple relationship types. Enterprises have relationships of various types with their partners. Some ways relationships can vary are: between major trading partners on the one hand and between minor trading partners on the other; between enterprises of roughly equal influence over the supply chain and between enterprises of unequal influence over the supply chain; and between enterprises with a high degree of technological sophistication on the one hand and between enterprises with an unequal degree of technological sophistication on the other. As should be understood, these different relationship types should be handled differently.

The present global collaboration manager can model enterprise relationships as a hub and spoke network, as described above and shown in FIGURE 4. In this embodiment, the four types of relationships are: Hub-to-Web; Hub-to-Van-EDI; Hub-to-Spoke and Hub-to-Hub. Each relationship-type has its appropriate usage.

With respect to Hub-to-Web, when people speak of E-Commerce today, they often imply an architecture where a web browser talks to some centralized server. This architecture has some advantages: the infrastructure to support this architecture is typically already in place; and all administration can be centralized on the server side. However, this architecture also has a big disadvantage in that it requires the presence of a human on the web-browser side. Hence system-to-system automation is not possible. Based on these and other pros and cons, this relationship type can be appropriate when an enterprise

needs to exchange data with either a minor partner or a partner with less technological sophistication.

5 With respect to Hub-to-VAN-EDI, the vast majority of electronic inter-enterprise commerce takes place today by sending EDI over Value Added Networks. The advantage of this approach can be that system-to-system integration is possible and it is currently supported today. Disadvantages of this approach are: large costs to send data over proprietary VAN's; high administrative costs
10 because of lack of true standardization; requirement for third party tools just to convert from the true "standard" to a form appropriate for the enterprise; no support for system-to-human integration; and no support for proprietary standards or corporate standards. Based on these and other pros and cons, this relationship type can be appropriate
15 when supporting a legacy VAN-EDI environment.

With respect to hub-to-spoke, this relationship type also enables system-to-system integration like VAN-EDI. Architecturally hub-to-spoke is a collaboration between a
20 hub engine and a spoke engine. The hub-to-spoke relationship can have advantages vis-a-vis VAN-EDI: it can use the public Internet to reduce network costs; administrative costs are much lower than VAN-EDI because a large portion of the hub-to-spoke relationship
25 infrastructure can be centrally deployed and administered; true objects (in addition to just data) can be exchanged allowing for much more advanced collaborations; and multiple semantic standards can be supported including EDI, I2-CDM and Proprietary Community Standards. Based on the
30 characteristics above, the hub-to-spoke relationship can be appropriate between enterprises that wish to perform sophisticated system-to-system collaboration. It can also

be appropriate where no I2 TECHNOLOGIES' software is present in either of the enterprises. This is because the hub-to-spoke relationship can be centrally deployed by the hub enterprise.

5 With respect to hub-to-hub, the relationship is similar to hub-to-spoke except that it takes place between two hub engines rather than a hub and a spoke engine. Based on this characteristic, the hub-to-hub relationship can be appropriate between enterprises that wish to perform
10 sophisticated system-to-system collaboration. Further, the hub-to-hub relationship can be appropriate when two enterprises have individually and separately purchased RHYTHM-GCM and have set up hub engines.

15 There are differences between hub engines and spoke engines. In general, a hub engine's capabilities are a superset of a spoke engine's capabilities. The following table provides an example of some of the differences.

TABLE 1		
	Spoke Engine	Hub Engine
20 Purchasing and Deployment	Spoke engines are bundled with a hub engine. Hence a hub enterprise will typically purchase a hub engine and a number of spoke engines which it can deploy out to its partners.	Sold separately.

	Relationship types supported	Can only support the hub-to-spoke relationship. Additionally, each spoke engine can only communicate with a particular hub engine (its owning hub).	Supports hub-to-hub, hub-to-spoke, hub-to-web and hub-to-VAN-EDI relationship types.
	Authoring Collaborations	Can view but not author a collaboration.	Can view and author a collaboration.
5	Internal-User Roles.	Supports a single internal-user role.	Supports multiple internal- user roles.

Security

10 A further problem with collaboration is the challenge
 of providing comprehensive security.
 Before enterprises can collaborate effectively, the
 security issue needs to be addressed. There are many
 different facets to security in a collaborative context.
 Any multi-enterprise collaborative framework should address
 15 all of these different facets. The requirements for a
 collaborative security framework can include that: data
 exchanged between two partners should only be seen by the
 two partners; data exchanged between two partners should be
 tamper-proof; an enterprise should be able to verify that
 20 a partner is who it claims to be; the framework should not
 introduce new security holes into a partners' network; and
 the framework should be relatively easy to set up and
 administer.

25 A secure collaborative framework can be provided by
 implementing a comprehensive security strategy to address
 the above requirements. In one embodiment, the strategy

has three different aspects to it: technological security, a permissibility framework and data partitioning.

Technological security can refer to the technological means used to guarantee security. This security can be used to provide: privacy, authentication and data integrity. Privacy ensures that no unauthorized person can see the data. Authentication involves authenticating that the parties in the collaboration are really who they claim to be. Data Integrity involves making it impossible for an unauthorized person to modify data being sent in any fashion.

The precise security approach can vary based on the relationship type described earlier. For example, one scheme is detailed in the table below:

TABLE 2		
Relationship Type	Technological Approach	Provided By
Hub-to-web	HTTP-over-SSL 3.0 (e.g., Diffie-Helman)	Global Collab Workspace
	HTTP-over-SSL 3.0 (e.g, RSA)	
	IIOP-over-SSL 3.0	Global Collab Workspace
Hub-to-spoke	HTTP-over-SSL 3.0 (e.g, Diffie-Helman)	Global Collab Workspace
	HTTP-over-SSL 3.0 (e.g., RSA)	Global Collab Workspace
	IIOP-over-SSL 3.0	Global Collab Workspace
Hub-to-hub	TCP/IP-over-SSL 3.0	Global Message Bus

	Content-based Encryption	Global Message Bus
Hub-to-VAN EDI	Security handled by VAN.	VAN

5 As can be seen from the table, all of the relationship types, with the exception of Hub-to-VAN EDI, could support security via SSL 3.0.

10 SSL 3.0 is an industry standard protocol used to support public key encryption over a socket-based connection and provides: privacy, client as well as server authentication, data integrity and certificate management. SSL 3.0 is a higher level protocol into which several public-key cryptography algorithms can be plugged including RSA and Diffie-Helman.

15 Once the SSL handshake is complete, the next step is username-password authentication. This provides authentication beyond what SSL 3.0 itself provides. Passwords can be stored using PKCS5 password-based encryption (an RSA standard). Once a user or spoke is authenticated, it is returned an Access Token. This access
20 token has an administrator-specifiable lifetime. A user can then access the system for the duration of validity of the access token. This has the beneficial effect of not requiring authentication on each access. Each application which is accessed, authenticates the access token by
25 validating the signature (which is a digest encrypted using the Security Manager's private key) of the Security Manager.

30 The technological security framework is a portion of the security scheme. The other portion has to do with the design of the collaborations themselves. The framework should allow enterprises to easily attach permissibilities

to various actions that other enterprises can perform on it. The global collaboration workspace can support a hierarchical permissibility model with individual permissibilities attached to different data elements in the hierarchy. In particular, it can support user-specific and spoke-specific read, write, take and subscribe permissibilities. Hence, enterprises can finely tune who can read what data, who can write what data, who can take what data and who can subscribe to write-notifications on what data.

A third element in the collaboration framework security strategy is the ability to partition data across various collaborative workspaces. In particular, the collaborative workspaces are split into an internal collaborative workspace and an external collaborative workspace. Only data that needs to be truly shared with partners is in the external collaborative workspace. The rest is in the internal collaborative workspace. The external collaborative workspace is designed to sit either outside the corporate firewall or in an Extranet or DMZ. The collaboration framework design does not require the external collaborative workspace to make connections through the corporate firewall into the Intranet (although it could).

In one embodiment, global collaborations can use both the external and internal collaborative workspaces. Local collaborations can use only the internal collaborative workspace and are hence completely invisible to partner enterprises. Even for global collaborations only the relevant portions use the external collaborative workspace. Furthermore, because of the permissibility framework

described earlier, each partner enterprise can only see (read, write, take, subscribe) to its own data.

FIGURE 8 is a block diagram of one embodiment of a security configuration for a hub-to-spoke and hub-to-web case. As shown, a hub enterprise 50 is coupled to and communicates with an internal global collaboration workspace 52 and an external global collaboration workspace 54. A spoke enterprise 56 and a web enterprise 58 connect through a web server 60 to the external global collaboration workspace 54. Spoke enterprise 56, like hub enterprise 50, has an internal global collaboration workspace 62. The enterprises 50, 56 and 58 can be protected by associated firewalls, while the extranet formed by web server 60 and external global collaboration workspace 54 can be protected by a filtering router and communication via HTTP over SSL 3.0.

FIGURE 9 is a block diagram of one embodiment of a security configuration for a hub-to-hub case. As shown, a hub enterprise 64 and a hub enterprise 66 can communicate across an SSL 3.0 protected TCP/IP connection. The communication can be between separate global message brokers 68 and 69. Both hub enterprises 64 and 66 are protected by a firewall, as shown.

Inter-Enterprise Workflows

One of the problems with multi-enterprise decision support can be that there is no closed loop collaboration. Instead, data may be lobbed from one enterprise to the next with no coherent workflow. In order to implement closed loop collaboration, support for creating multi-enterprise workflows is necessary. The present global collaboration manager and designer can make it possible to construct,

deploy, monitor and change sophisticated multi-enterprise workflows.

5 In general, a "workflow" can be a set of "activities" joined together by data flows that together accomplish some task. Workflows are typically executed on workflow engines. A "distributed workflow" can refer to a workflow that is executed on multiple workflow engines. In other words, different portions of the workflow execute on different engines. A "node" can refer the abstract
10 entities on which different workflow engines of a distributed workflow run, and a "node group" can be a set of nodes grouped by some characteristic. A "multi-enterprise distributed workflow" can be distributed workflows where the nodes are enterprises.

15 Parameterization of workflows can be important for enterprise collaboration. A "parametric workflow" is a workflow that is parameterized over some variable and can be regular or distributed. Instantiating the parametric workflow with different values of the parameter variable(s) produces different instances of the workflow. A
20 "distributed workflow parameterized over nodes in a node group" can refer to distributed workflows where the parameters of the workflow are the nodes in a node group. Hence, when the workflow is instantiated it is tailored to
25 a particular node in a node group.

There are several important features to the workflows that can be supported by the present global collaboration. These workflows can be strongly typed. Strong typing can be essential in producing robust, error-free workflows. In
30 essence, strong typing guarantees the type of a message at design time. For example, if the workflow is designed to send a Bill of Materials, then strong typing ensures that

it is physically impossible that an object other than a Bill of Material is sent. For a workflow designed with the global collaboration designer and executed with the global collaboration manager, it can be made impossible to even
5 send an object of an incorrect type. This capability is important to producing robust, error-free workflows.

Despite strong typing, there are, for example, two scenarios in which wrong object types could conceivably be passed in the workflow: due to an error on the workflow
10 designer's part; and a malicious attempt by someone to undermine the workflow. Both of these scenarios can be handled. The first can be handled by making it impossible for an error in design to lead to such a scenario. The second can be handled by making the data flows tamper-proof
15 by using public key cryptography or other encryption scheme (integrity characteristic) as described above.

Another important feature is support for workflows parameterized over groups. Some multi-enterprise workflows involve a large number of enterprises. In such cases it
20 can become impractical to create individualized workflows for each partner. Instead it can be advantageous to create workflows that are parameterized over groups of partners. For example, in the realm of procurement, two groups may be primary suppliers and secondary suppliers. The primary
25 suppliers group could have one type of workflow, and the secondary suppliers group could have another type of workflow. Group-based workflows can be parametric in the sense that, at run time, an actual workflow can be created specific to a member of a group.

30 In the multi-enterprise context, an enterprise may collaborate, for example, with potentially hundreds or thousands of other enterprises. Each collaboration or

multi-enterprise workflow can be potentially (and typically) unique. However, designing thousands of specialized workflows with an enterprises' partners is neither desirable nor feasible. On the other hand, many of these workflows are simply parametric variations on an underlying parameterized workflow. For example, a company A may be collaborating (on sales) with retailers, distributors, direct sales etc. Hence, it makes sense to group the various partners. An example grouping may be: WalMart; Sears; Rest of Retailers besides WalMart and Sears (group); Primary Distributors (group) and Secondary Distributors (group). Now, the workflows with all the members, for example, of the primary distributors group are variations on an underlying parametric distributed workflow, parameterized over the particular distributor in that group.

Workflows parameterized over groups can be supported by a HETEROCASTING workflow definition technique. The HETEROCASTING definition technique generally involves using a parameterized workflow definition to instantiate heterogeneous workflows based upon differences in the parameters. Thus, the HETEROCASTING definition technique allows a non-parametric distributed workflow to be easily (through a visual design tool) be made parametric over nodes in a node group. There can be two primary workflow activities used to accomplish this definition: a HETEROCAST split activity and HETEROCAST join activity. All activities between a HETEROCAST split and a HETEROCAST join are parameterized over the nodes of a node group that these activities correspond to.

FIGURE 10 is a diagram of one embodiment of designing an inter-enterprise workflow that includes parameterization

over groups. As shown, the workflow can begin with a listening activity 70 that waits for some event. Activity 70 can be linked to parallel activities 71 that links to a sub-workflow 72 and to a heterocast split 73. Sub-workflow, itself, can include a workflow definition. With respect to HETEROCASTING, the workflow after heterocast split 73 then becomes parameterized. Thus, in the example of FIGURE 10, activity 74 is a parameterized activity. After activity 74, a heterocast join 75 receives flow from activity 74. Sub-workflow 72 and heterocast join 75 are linked to a synchronous or asynchronous join 76 which, in turn, links to an integrated event 77 (e.g., multicasting). A workflow like that of FIGURE 10 can be designed using the present global collaboration designer and can allow full representation of workflow for inter-enterprise decision support. This workflow can then be instantiated and implemented through the present global collaboration manager.

FIGURE 11 is a diagram of one embodiment of managing change by modifying a design of a workflow. As shown, an initial workflow design can have an event 70 linked to a parallel activity split 71. Between activity split 71 and a join 76, there can be, for example, two activities 78. This workflow, once designed, can be instantiated and implemented using the global collaboration manager. If a change needs to be made to the workflow, the global collaboration designer greatly alleviates the trouble of making the change. For example, a new activity 79 can be added between split 71 and join 76. The workflow can then be centrally reinstantiated and implemented.

In particular, the HETEROCAST technique can allow the construction of distributed workflows parameterized over

nodes in a node group. This can allow a huge productivity gain over designing individual workflows for individual group members. Further, this technique makes rapid design and prototyping of sophisticated inter-enterprise workflows with hundreds or thousands of partners feasible. The technique should be distinguished from conventional "multicasting" in which identical messages are sent out to the various nodes (partners). In essence, multicasting allows you to design a single workflow that runs identically across multiple nodes. This differs from the HETEROCASTING technique, where the workflows run differently based on which node they are running across.

A third important feature is support for role-based workflows. Role-based workflows allow workflows to be specified using generic roles. This capability allows the creation of generic or templated workflows that can be instantiated in various scenarios. For example, the role types can be: partner roles, spoke roles; spoke group roles; web roles; web group roles; user roles. As an example of roles, partner roles refer to the different roles played by partners. Thus, one partner role in the case of procurement is primary supplier and secondary supplier.

Role-based workflows can lead to the concept of three different phases in the design and execution of a workflow. The design phase is the phase in which role-based workflows are defined. The instantiation phase is the phase in which roles are mapped to instances. For example, primary supplier may be mapped to a first company, and PO_approver may be mapped to John Doe. Third, the run time phase can be the phase in which the instantiated workflow runs.

A further important feature is the integration of automated workflows with user-oriented workflows. Workflows can often be described as having two varieties: automated system-to-system workflows, and user interface workflows. While there are workflows that are completely automated, and there are workflows that are completely user driven, most workflows have automated as well as user interface elements. The present global collaboration manager and designer do not need to make this artificial distinction between workflow types. Hence, the workflows can be automated in parts and interact with users in other parts. Both the automated parts and user parts can span multiple enterprises.

Integration with Outside World

FIGURE 12 is a diagram of one embodiment of integration of a workflow with the outside world. As described in the previous section, sophisticated inter- and intra-enterprise workflows can be created. These workflows can be composed of activities strung together in various configurations. There is no restriction on what the different activities of the workflow can do, yet one of the major tasks of these activities is to integrate with the outside world. FIGURE 12 shows how a workflow can be integrated with the outside world using a component-based approach to integration. The components can include accessors 80, transformations 82, transfer objects 84, adaptors and flows 86.

The global collaboration manager can support a component-based integration model. The component-based integration model allows flexibility in structuring the integration. There can be two types of components:

primitive components and compound components. Primitive components can include accessors 80, transformers 82 and transfer objects 84. Compound components include adaptors and flows 86. Compound components are built in terms of primitive components. In this scheme, accessors 80 are used to access an external source such as SCP (SUPPLY CHAIN PLANNER), SAP, a relational database, web servers, email, message buses etc. Accessors 80 can be used to read, write or listen to sources and destinations of data. Transformers 82 can be used to transform data from one form to another form. Transfer Objects 84 are objects that can be passed from activity to activity or from enterprise to enterprise. Transfer objects 84 can be optionally convertible to EDI, XML, CORBA structures etc. Accessors 80 and Transformers 82 can be strung together to form flows. An entire flow can be executed in a single activity as shown in FIGURE 13.

FIGURE 13 is a diagram of one embodiment of a data flow running in a single activity 92. As shown, a data source 90 can be accessible from and provide data to an accessor component 94. Accessor component 94 then can pass data through transformer components 96 and 98 which provide data to a second accessor component 100. Data can then be stored in a data destination 102.

FIGURE 14 is a diagram of one embodiment of a data flow split across multiple activities 104 and 106. As shown, the flow of FIGURE 14 differs from that of FIGURE 13 in that transformer components 96 and 98 are within separate activities 104 and 106 and communicate by a transfer object. Multi-enterprise data flows can be based on the model of FIGURE 14 rather than that of FIGURE 13.

With respect to transformations, in one embodiment, two fundamental transformation types can be supported: I2-CDM based transformations and direct transformations. I2-CDM based transformations are based on I2 TECHNOLOGIES' COMMON DATA MODEL (CDM). The CDM is an abstract schema that is available in both relational and object forms.

FIGURE 15 is a block diagram of one embodiment of an I2-CDM based transformation model. As shown, transformers and accessors can be coupled to transform a application data into a CDM data object 110 and vice versa. For example, a SUPPLY CHAIN PLANNER (SCP) object 112 can be created by an SCP accessor from SCP data 114. SCP object 112 can then be transformed by an SCP-CDM transformer into a CDM object 110. Analogously, an SAP object 116 can be created by an SAP accessor from SAP data 118. SAP object 116 can then be transformed by an SAP-CDM transformer into a CDM object 110. The SAP accessor and transformer, as with other accessors and transformers, can be combined into a standard SAP-CDM adapter 120 that can be used for CDM-based transformations other components. As another example, a BAAN object 122 can be created by a BAAN accessor from BAAN data 124. BAAN object 122 can then be transformed into a CDM object 110 by a BAAN-CDM transformer. These transforms work in the other direction as well.

FIGURE 16 is a diagram of one embodiment of a direct transformation. In direct transformers, objects are converted from one form to another without passing through an intermediate format. For example, as shown in FIGURE 16, SUPPLY CHAIN PLANNER (SCP) data 130 can be accessed by an SCP accessor to create an SCP object 132. SCP object 132 can then be directly transformed to a FACTORY PLANNER (FP)

object 134. FP object 134 can then become FP data 136 through an FP accessor. This data flow can operate in the other direction as well.

5 In these processes, there are various levels of granularity at which access and transformation can take place including the relational (table), generic object (tree, graph, matrix etc.) and specific object (Bill of Material, Plan etc.) levels. Sometimes access may only be available at one level (say tables), but transformation may
10 be more appropriate at another level (say generic object). For example, hierarchical aggregation (a form of transformation) is often appropriate on a tree object. However, the data may only be accessible in a tabular form. In this case, for example, the data should be accessed at
15 the tabular level, transformed into a tree, and then have the hierarchical aggregation applied to it.

FIGURE 17 is a diagram of one embodiment of different access and transformation levels. As shown, access and transformation can have three levels. A first level 140
20 can involve table access and transforms. A second level 142 can involve generic object (tree, graph, etc.) access and transforms, and a third level can involve specific object (build-of-materials, plan, etc.) access and transforms. In addition to transforms between
25 application formats, there can also be transforms between the three levels, as shown.

Deployment of Collaborations

30 One important factor in a multi-enterprise collaboration system is the ease with which the collaboration can be deployed. As discussed, the present global collaboration manager can support four different

kinds of partner relationships: hub-to-web, hub-to-spoke, hub-to-hub and hub-to-VAN-EDI. Of these four, hub-to-web has all the deployability characteristics of traditional web applications. Hub-to-VAN EDI can be deployable to the extent that it leverages an existing VAN-EDI infrastructure. While the hub-to-web relationship is highly deployable, it can suffer from the problem of requiring a human on the web side of the relationship. In other words, it may not be suited to system-to-system collaboration.

The hub-to-spoke solution can provide maximal deployability in the system-to-system collaboration environment. In the hub-to-spoke realm, the spoke engine is analogous to the web browser, and the spoke portion of the collaboration is analogous to a web page or applet. Similar to a web-page or applet, the spoke portion of the collaboration can be centrally designed and deployed to the remote spoke engines. Unlike a web-page or applet, there may still be integration that needs to be done remotely. This remote integration may be unavoidable but can be circumscribed and precisely defined by the spoke portion of the collaboration.

Another aspect of deployability is handling versioning. Collaborations once designed and deployed are likely to need changing (in various different ways) as time progresses. It can be important that subsequent versions of collaborations be as easily deployable as initial versions. The present global collaboration manager can provide complete support for versioning and centralized redeployment of collaborations. Further, different versions of collaborations can be run simultaneously without impacting each other. This allows an existing

version to be gracefully phased out while another version is phased in.

Another element of the deployability of the present global collaboration manager is the leverage of existing infrastructure. This element is evident, for example, in the support of the hub-to-spoke relationship over existing web protocols. Supporting hub-to-spoke over existing web protocols can be important to rapid deployment since it does not require modification or reconfiguration of an existing web infrastructure. A large time savings in this regard can come from not having to modify carefully designed firewall and security infrastructures that may already be in place.

Supporting Many-To-Many Collaborations

The present hub-and-spoke architecture provides easy manageability and deployment. However, in practice enterprises collaborate with many enterprises which in turn collaborate with still other enterprises. Hence, enterprises often form a collaborating web or graph. This can be supported via the ability to substitute a hub engine for a spoke engine at any time. This substitution ability allows many-to-many collaboration webs to be grown organically rather than all at once.

FIGURE 18 is a diagram of one embodiment of substituting a hub engine for a spoke engine within a collaboration. As shown, an enterprise (E1) may deploy a hub engine 150 on itself and a spoke engine 152 at all of its partner sites. In particular, a spoke engine 154 may be at a partner site (E2). If the partner site (E2) wishes to design and control its own collaborations, it can replace spoke engine 154 with a hub engine 156. From E1's

perspective, E2 can still be a spoke in E1's collaboration. However, this spoke now runs on a hub engine 156 which can control its own collaborations with spoke engines 158. Further, spoke engines 160 and 162 might be associated with
5 a third entity (E3) that interacts with both hub engine 150 and hub engine 156 on behalf of E3.

Extension of Framework

10 An important aspect of the present framework is extensibility. Without extensibility, the framework may not be able to handle new situations and challenges with which it is confronted. There can be several different dimensions to this extensibility. For example, one primary
15 area of extensibility is in the area of semantic object standards. If supported standards do not suffice for a particular problem, then the framework can be augmented with new semantic standards. Additionally the framework allows the building of proprietary semantic standards.
20 Further, the framework can be extended by adding new accessors, transformers, adapters, etc. The standard component library can be extended both generally and by end-users.

Object Workspace

25 FIGURE 19 is a block diagram of one embodiment of a computer workspace 200 in a computer system 205. Computer workspace 200 includes a plurality of memory slots 210 in communication with a permissibility framework manager 220 and an event manager 230. Permissibility framework manager
30 and event manager 230 may or may not be resident within computer work space 200. Computer workspace 200 is accessed by network nodes 240 via network 250. Generally,

permissibility framework manager 220 controls access to memory slots 210 within computer workspace 200 by network nodes 240. Event manager 230 generates messages to network clients 240 in response to events or conditions associated with data or objects that are stored in memory slots 210.

Network 250 comprises any combination or number of hubs, routers, bridges, gateways, switches, or any other association of suitable communication devices and related software that transmits data between network nodes 240.

In one embodiment, network 250 comprises a network implemented independently or in connection with a wide area network (WAN) or a local area network (LAN), such as an Ethernet network, a token ring network, or a fiber distributed data interface (FDDI) network. Network 250 supports higher level connectionless protocols such as Internet Protocol (IP), higher level connection-oriented protocols such as Frame Relay, or any other suitable networking protocol. Network 250 may be used in a multi-enterprise collaboration to implement workflows including activities that take place among or between more than one enterprise. Each network node 240 of network 250 in a multi-enterprise collaboration may be associated with a different enterprise, allowing communication and coordinated performance of activities and workflows between enterprises.

Network nodes 240 may be any terminal, server, client, hub, spoke or other device connected to network 250. Each network node 240 is associated with a particular enterprise. Network nodes 240 may or may not participate in a particular workflow or process. Network nodes 240 may access workspace 200 as part of a workflow or collaboration

or as part of other activities or processes not associated with a workflow or collaboration.

Each memory slot 210 stores data and objects. As used herein, each means each of at least a subset of the identified items. Objects can be Java objects, C++ objects, Corba objects, or other structures that are capable of storing both information and behavior. Memory slots 210 may be any memory structure, whether queued or randomly accessed, capable of holding data or objects. Memory slots 210 may include, for example, any data structure, hashtable or memory array. Memory slots 210 may, for example, contain a plurality of objects enqueued within memory slots 210 or may contain only a single object. Memory slots 210 may be stored to a disk or other storage medium or may be maintained in memory during any processes or activities conducted by computer system 205. Memory slots 210 maintained in memory may be stored on random access memory, for example, enhancing the speed at which their contents are accessed over those memory slots 210 stored to a disk or peripheral component. Memory slots 210 stored to a disk or other storage media are accessible at a lower speed than memory slots 210 stored in memory but are nonvolatile allowing for the storage of persistent data or objects. Memory slots 210 can be arranged in an organizational hierarchy defined by a programmer, user or other suitable mechanism. Such a hierarchy allows easy categorization of and reference to memory slots 210 as described below in reference to FIGURE 20.

Permissibility framework manager 220 maintains and controls access to memory slots 210. Permissibility framework manager 220 may include any combination of hardware and software capable of maintaining access rights

to memory slots 210 and controlling access to memory slots 210 based on the access rights maintained. In one embodiment, access rights include the right of a node to: read from the contents of each of memory slots 210, write to the contents of each of memory slots 210, remove any of the contents to each of memory slots 210, and subscribe to and unsubscribe from event notification for one or more specified memory slots 210.

Event manager 230 generates messages to nodes 240 in response to a particular event within memory slots 210. Event manager 230 may include any combination of hardware and software capable of generating messages and initiating the routing of such messages to a particular network node 240 or other suitable element of the computer system 205. Network nodes 240 that have access rights to subscribe to event notification for a particular memory slot 210, as determined by permissibility framework 220, and that exercise such subscription to the particular memory slot 210, are notified by a message generated by event manager 230 each time an event requiring notification occurs that is related to the particular memory slot 210.

For example, one of network nodes 240 may have access rights to subscribe to notification for a particular memory slot 210, as verified by permissibility framework manager 220, and may subscribe to notification for the particular memory slot 210 each time that the particular memory slot 210 is written to by one of network nodes 240. Until the one network node 240 unsubscribes from notification for that particular slot, each time that the particular memory slot 210 is written to by one of network nodes 240, event manager 230 will generate a message and initiate the routing of such message to the one network node 240

indicating that the particular memory slot 210 has been written to.

5 In one embodiment, access rights to subscribing and unsubscribing to event notification for a particular memory slot 210 may vary based on the classification of an event. For example, a particular network node 240 may have subscription rights to be notified whenever a particular memory slot 210 has data removed but not when the same memory slot 210 is written to, or the network node 240 may have subscription rights to both, or subscription may encompass notification in response to either data being removed or written in such memory slot 240. Access rights granted by permissibility framework 220 and subscription to messages from event manager 230 may also be grouped in any suitable combination. For example, event notification may be subscribed to for an individual memory slot 210, all of memory slots 210, or a subset of memory slots 210 selected based upon indicated criteria. Likewise, access rights can be granted by permissibility framework 220 to all of memory slots 210 or any subset thereof.

20 FIGURE 20 illustrates one embodiment of the workspace illustrated in FIGURE 19. Workspace 300 is a workspace for storing data and objects in memory slots 310 that are arranged in a hierarchical framework. The exact nature of the hierarchy and the placement of particular memory slots within the framework is definable by a workspace administrator, other user, or other suitable mechanism.

25 Generally, workspace 300 is organized into sections 305 that are further separated into individual memory slots 310 and groups 320. Groups 320 may in turn contain memory slots 310 and/or be separated into further subgroups. In FIGURE 20, memory slots 310 are designated by section

number, group number if applicable, and then designated by s and an identification number within the applicable section 305 or group 320. For example, the memory slot 310 identified by the nomenclature "section 2.group1.s2" denotes a slot named "s2" in group 1 of section 2 of the workspace. Other hierarchy or organizational scheme may be substituted for the hierarchy described herein. A hierarchy such as the one described allows easy categorization and grouping of memory slots 310. Such categorization can be used to easily group memory slots 310 for purposes of maintaining access rights. For example, a network node 240 from FIGURE 20 may only be granted a particular access right to a specific tier of memory slots 310. Such a tier may be all memory slots 310 at the group level, for example, or memory slots 310 designated by section1.s1 and section section2.s1 in FIGURE 20. Access rights may also be granted to a particular section 305, group 320, subgroup, or combination thereof.

FIGURE 21 is a block diagram of a computer system 405 that uses at least one agent 460 initiated at one of a plurality of network nodes 440 to access an object workspace 400 across a network 450. Network nodes 440 are connected via network 450 to workspace 400. Generally in computer system 405, agent 460 is initiated at one of network nodes 440 to travel across network 450, access workspace 400, and perform operations or tasks in workspace 400 based on programmed commands.

Workspace 400 includes memory slots 410 and is embodied with any or all of the characteristics of workspaces 200 and 300 as described in reference to FIGURES 19 and 20. Network nodes 440 are associated with a particular enterprise and are embodied as described in

reference to network nodes 240 of FIGURE 19. Network 450 is a network connecting network nodes that are involved in a particular collaboration and is embodied as described in reference to network 250 of FIGURE 19.

5 Agent 460 can be any object-oriented agent or other item that possesses the characteristic of autonomy. Autonomy, as defined herein, is the ability to be programmed with one or more goals and to attempt to satisfy those goals independently of the originating program or
10 application, even when moved into a network and onto other platforms such that contact is lost with the originating program. Agents 460 are programmed with the goal of fulfilling collaboration activities. Collaboration activities may be activities involved in the performance of
15 a workflow or activities performed outside of a workflow that an enterprise or facility within a collaboration may initiate to perform administration, optimization or execution tasks related to data or objects that are in workspace 400. Collaboration activities may also be queries
20 intended to collect data or manipulate data or objects that are in workspace 400. An administrative agent 460, for example, may be initiated to modify characteristics of the memory slot hierarchy as described in reference to FIGURE 20. An administrative agent 460 may also be initiated to
25 change the subscription profile of an originating network node 440, thereby altering the type of event notification received by that network node 440. A specific agent 460 may even originate from a specific network node 440 that has the authority to alter the permissibility framework,
30 and thereby enable agent 460 to modify access rights of one or more network nodes 440 to specified memory slots 410. An optimization agent 460 may be initiated to modify

characteristics or behavior of an object stored in workspace 400 to improve that object's performance during workflow determination or execution. An execution agent 460 may perform a series of tasks or calculations using objects in workspace 400 and return a result of such tasks or calculations to its initiating network node 440. A query agent 460 may perform complex queries on workspace 400 and return a result of such queries to its initiating network node 440. Agents 460 can operate inside workspace 400 as well, returning results to another agent 460 or network node 440.

Workspace 400 may, like workspaces 200 and 300, have a permissibility framework and/or event manager capable of interacting with or responding to agent 460. For example, agent 460 may need to interact with a permissibility framework, as described in FIGURE 19, in order to verify that the originating network node 440 of agent 460 has access rights to a particular memory slot 410 that contains objects or data with which agent 460 wishes to interact. Similarly, a event manager may respond to an operation performed by agent 460 within a particular memory slot 460 by sending an event notification message to subscribing network nodes 440 as described in reference to FIGURE 19.

FIGURE 22 is a flowchart illustrating one embodiment of a method of using the agent of FIGURE 21 to access a workspace. In step 510, a node of the network initiates an agent programmed to execute at least one task in a remote workspace. In step 520, the agent is sent across the network to a workspace that is coupled to a second node of the network. The workspace has a plurality of memory slots, each operable to store at least one object, a permissibility framework, and a event manager. In step

530, the agent indicates to the permissibility framework a specific memory slot that the agent wishes to access. In step 540, the permissibility framework verifies that the node of the network that initiated the agent has access rights to the indicated memory slot. In step 550, the agent performs one or more operations as described in reference to FIGURE 21. In step 560, the agent returns the results of the operations, if any, to the initiating node. In step 570, a event manager sends an event to a node of the network that has subscribed to event notification that corresponds to an operation the agent has performed and to the memory slot accessed by the agent.

Although the present invention has been described in detail, it should be understood that various changes, substitutions and alterations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.

WHAT IS CLAIMED IS:

1. A computer system for remotely accessing data in a multi-enterprise collaboration, comprising:

5 a network node associated with a first enterprise, the network node being in communication with a workspace across the network, the workspace being associated with a second enterprise and having a plurality of stored objects; and
an agent generated at the network node, the agent being operable to access the workspace via the network, the
10 agent being further operable to manipulate at least one of the plurality of stored objects within the workspace to perform a collaboration activity.

2. The computer system of Claim 1, wherein the agent
15 is an administrative agent.

3. The computer system of Claim 1, wherein the workspace is organized according to a hierarchy and wherein the agent is an administrative agent operable to modify
20 characteristics of that hierarchy.

4. The computer system of Claim 1, wherein the agent is an execution agent.

25 5. The computer system of Claim 1, wherein the agent is an optimization agent.

6. The computer system of Claim 1, wherein the agent
30 is a query agent.

7. The computer system of Claim 1, further comprising the workspace, wherein the workspace includes a plurality of memory slots, each slot being operable to store at least one object.

5

8. The computer system of Claim 7, wherein the memory slots are accessible by a plurality of network nodes and wherein access to the memory slots by the network nodes is controlled by a permissibility framework.

10

9. The computer system of Claim 1, wherein the workspace is an in memory workspace.

The computer system of Claim 1, wherein the agent is an optimization agent.

15

10. The computer system of Claim 1, wherein the workspace is a persistent workspace.

20

11. The computer system of Claim 1, further comprising the workspace, wherein the workspace includes an event manager, the event manager being operable to notify the network node in response to changes to the workspace.

12. A method of remotely accessing data in an object workspace, comprising:

initiating an agent at a first node of a network;

5 sending the agent across the network to a workspace at
a second node of the network, the workspace having a
plurality of memory slots, each memory slot operable to
store at least one object, the workspace further having a
permissibility framework to control access to the memory
slots;

10 using the permissibility framework to verify that the
agent has access rights to the specific memory slot; and
performing an operation using the agent and at least
one object stored in the specific memory slot.

15 13. The method of Claim 12, further comprising
sending the results of the operation to the first node.

20 14. The method of Claim 12, further comprising
sending a message to a third node of the network in
response to performing the operation.

25 15. The method of Claim 12, further comprising
changing an organizational hierarchy of the workspace in
response to performing the operation.

16. The method of Claim 12, further comprising the
first node subscribing to event notification in response to
performing the operation.

30 17. The method of Claim 12, further comprising
altering the permissibility framework in response to
performing the operation.

18. The method of Claim 12, wherein performing the operation comprises performing administrative tasks in the workspace.

5 19. The method of Claim 12, wherein performing the operation comprises optimizing the performance of the object.

1/12

FIG. 1

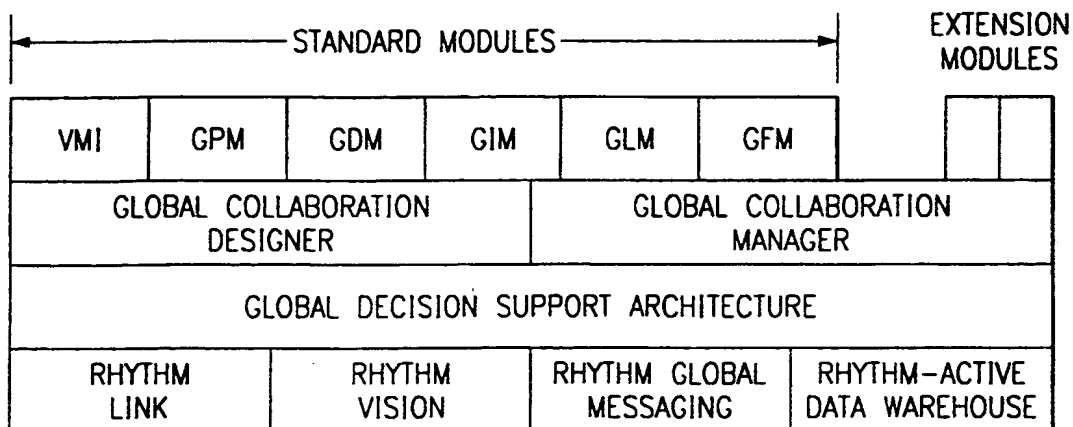
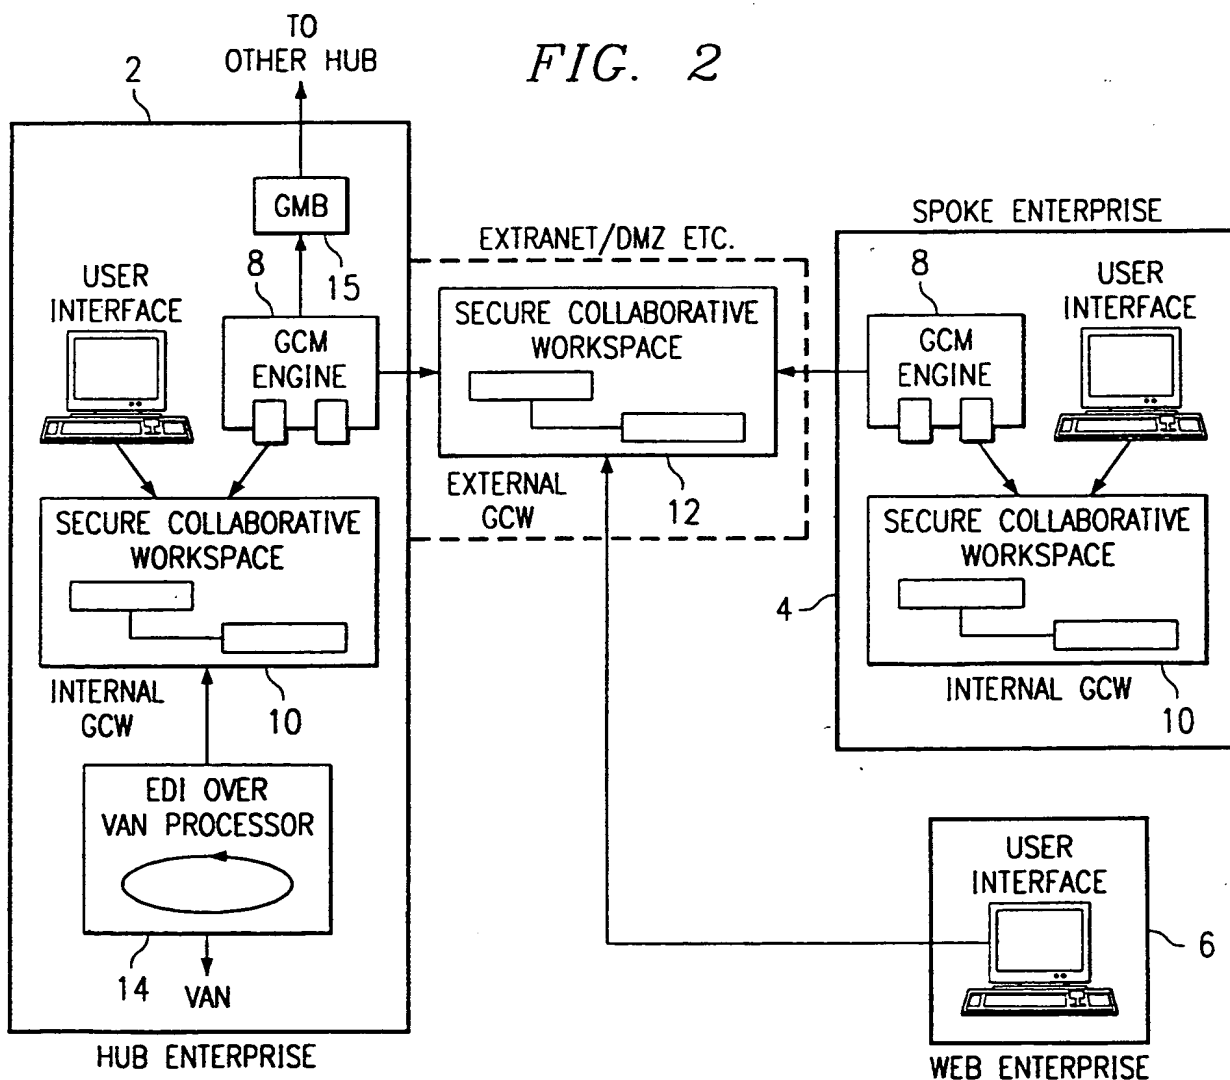

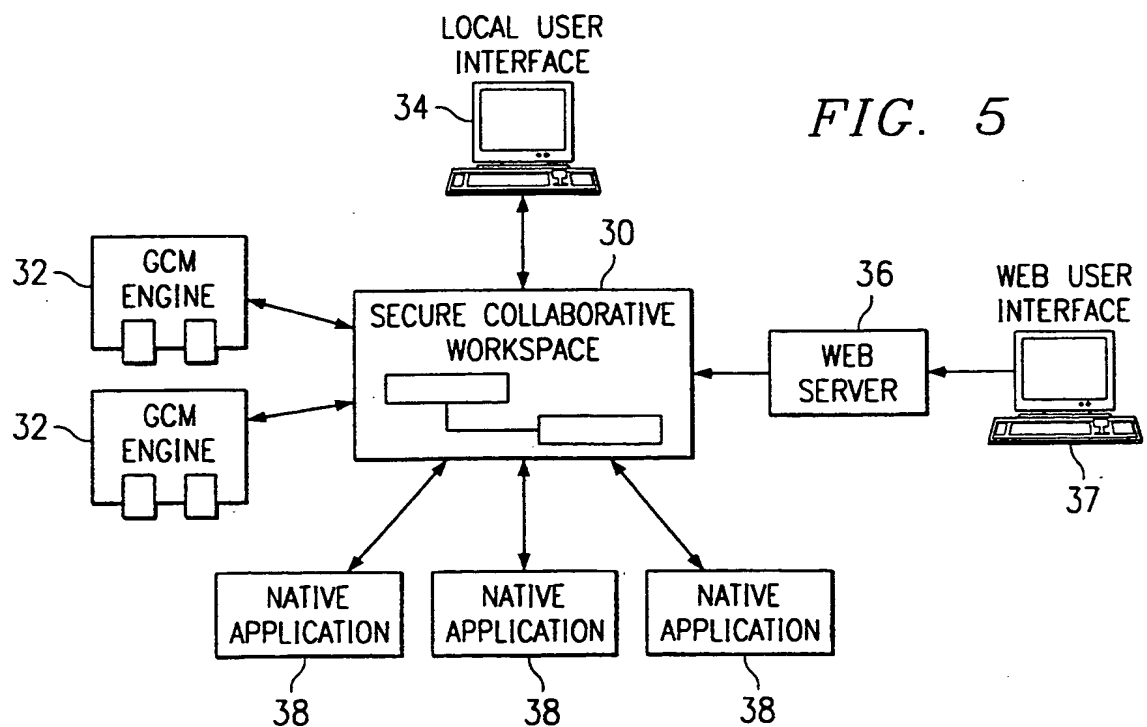



FIG. 2



2/12





4/12

FIG. 6

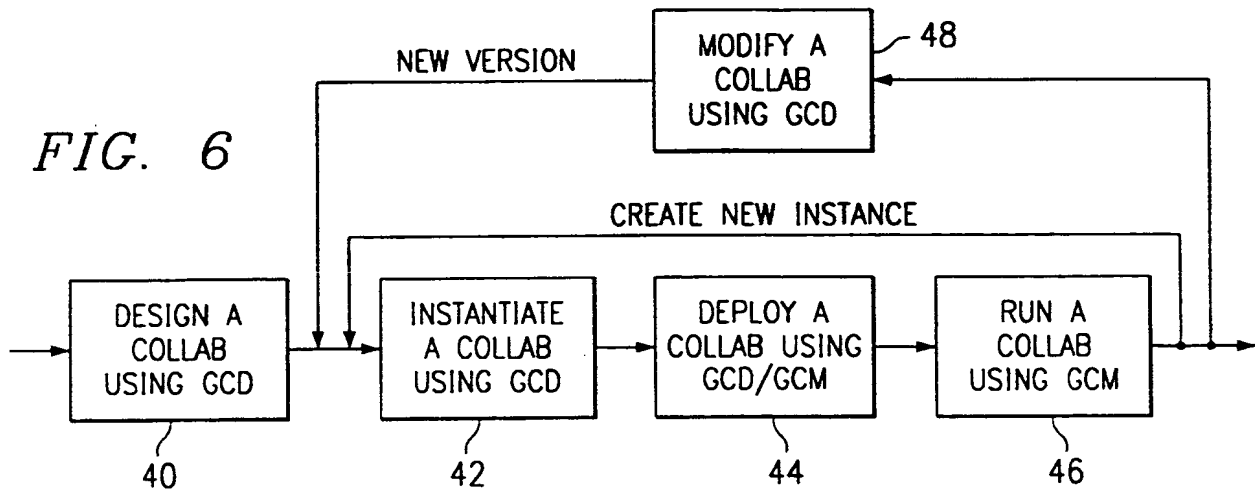


FIG. 7A




FIG. 7B






FIG. 9


5/12

FIG. 8





6/12


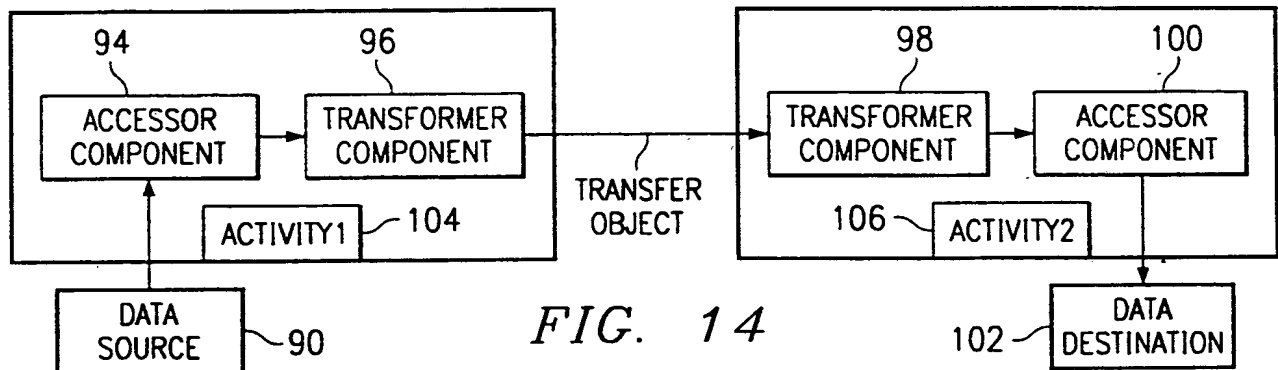

FIG. 10



7/12



8/12



9/12

FIG. 16




FIG. 18



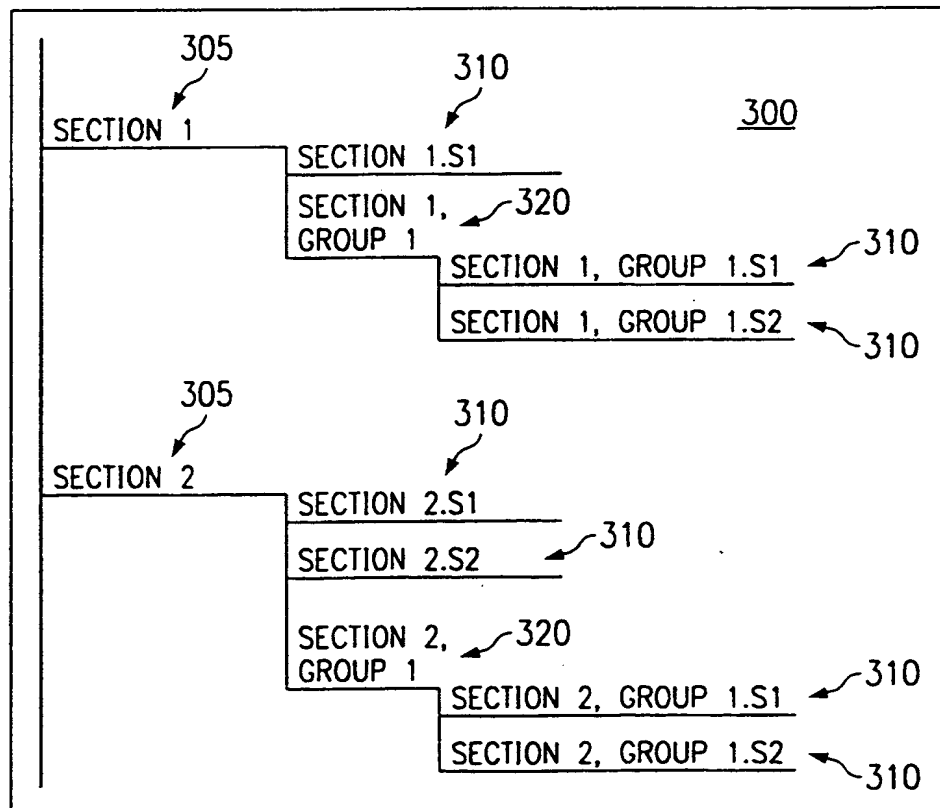
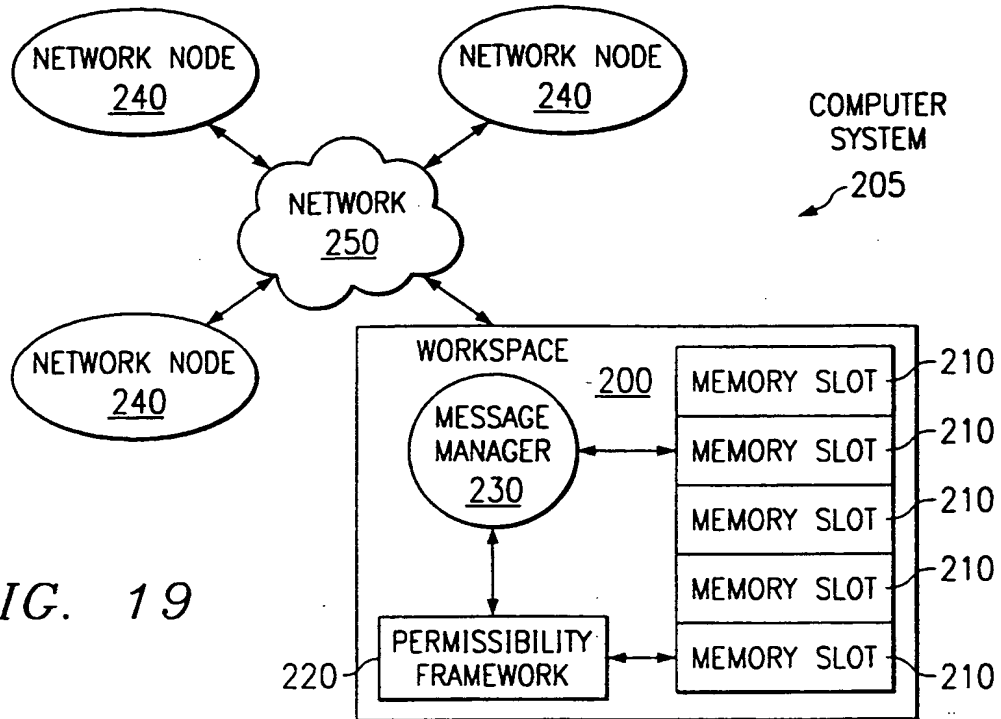


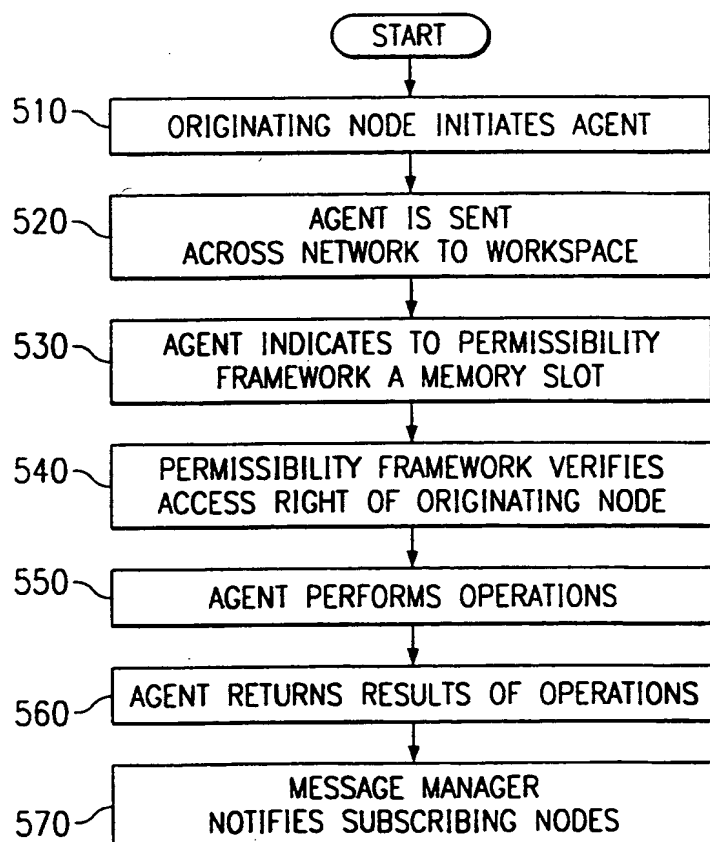
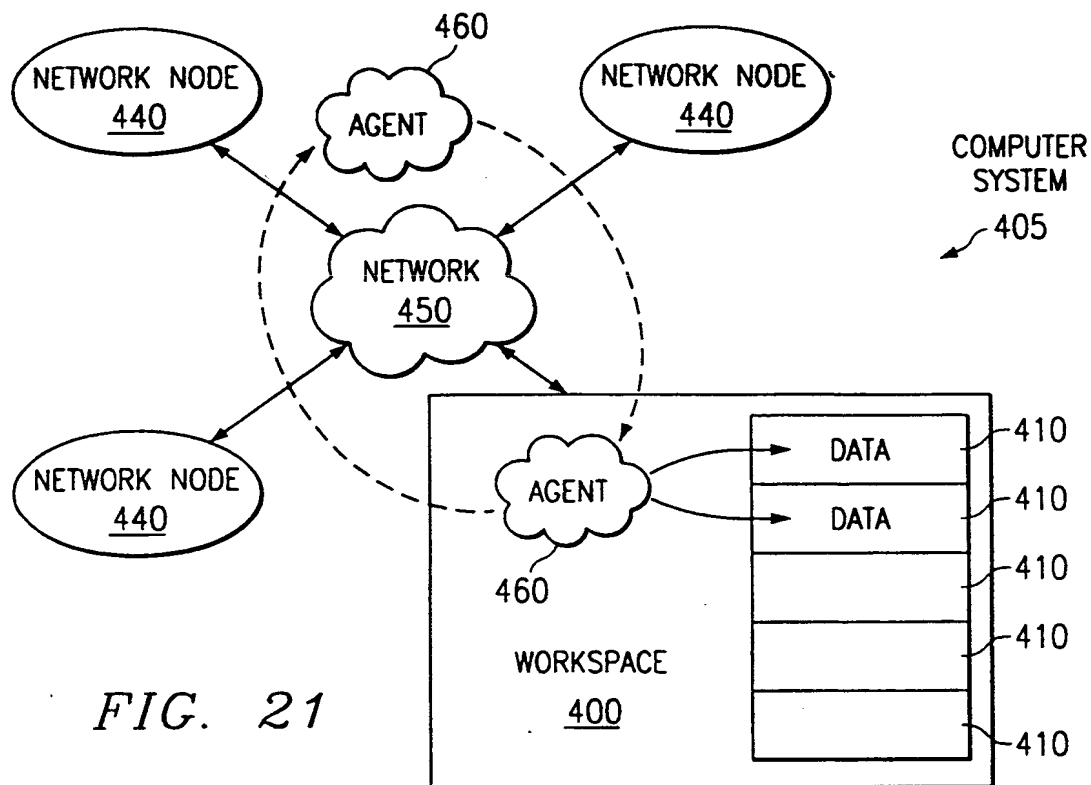
FIG. 17



11/12



12/12



INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/12347

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F17/60

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y A	WO 98 08177 A (I2 TECHNOLOGIES INC) 26 February 1998 (1998-02-26) abstract; claims 1-4; figures 1-5 page 2, line 1 -page 7, line 15 page 11, line 12 -page 14, line 5 ---	1-6 7-19
Y A	FOX ET AL: "The Integrated Supply Chain Management System" INTERNET CITATION, 7 December 1993 (1993-12-07), pages 1-11, XP002090173 page 2, line 23 -page 9, line 34 ---	1-6 7-19
	--- -/--	

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

20 October 1999

Date of mailing of the international search report

27/10/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Gardiner, A

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/US 99/12347

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 321 605 A (CHAPMAN WILLIAM ET AL) 14 June 1994 (1994-06-14) abstract; figures 1,2 column 3, line 10 -column 4, line 51 column 5, line 31 -column 10, line 18 ---	7-19
A	CIANCARINI P ET AL: "A multi-agent process centered environment integrated with the WWW" PROCEEDINGS - THE WORKSHOP ON ENABLING TECHNOLOGIES: INFRASTRUCTURE FOR COLLABORATIVE ENTERPRISES, 1 January 1997 (1997-01-01), pages 113-120, XP002090172 ISSN: 1080-1383 page 113, column 1, line 26 -page 114, column 2, line 10 page 116, column 1, line 1 -page 120, column 1, line 10 ---	1-19
A	US 5 023 773 A (BAUM RICHARD I ET AL) 11 June 1991 (1991-06-11) abstract; claims 12,13,20; figures 2,3 column 9, line 38 -column 12, line 37 -----	12-19

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/12347

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 9808177	A	26-02-1998	AU 3423597 A	26-02-1998
			CA 2195813 A	22-02-1997
			EP 0861474 A	02-09-1998
			GB 2316506 A,B	25-02-1998
			JP 10097574 A	14-04-1998
US 5321605	A	14-06-1994	NONE	
US 5023773	A	11-06-1991	CA 1313424 A	02-02-1993
			DE 3854616 D	30-11-1995
			DE 3854616 T	13-06-1996
			EP 0327707 A	16-08-1989
			JP 1228039 A	12-09-1989
			JP 1931903 C	12-05-1995
			JP 6061069 B	10-08-1994

THIS PAGE BLANK (USPTO)